



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/739,618	12/18/2000	John H. Howard	5181-59100	3682

7590 10/05/2005

LAWRENCE J. MERKEL
CONLEY, ROSE, & TAYON, P.C.
P.O. BOX 398
AUSTIN,, TX 78767-0398

EXAMINER

DUONG, THOMAS

ART UNIT PAPER NUMBER

2145

DATE MAILED: 10/05/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/739,618

Applicant(s)

HOWARD, JOHN H.

Examiner

Thomas Duong

Art Unit

2145

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 18 July 2005.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 2-10, 12-20, 22-27, 29-35, and 40-44 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 2-10, 12-20, 22-27, 29-35, and 40-44 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____.

- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____.

DETAILED ACTION

Response to Amendment

1. This office action is in response to the applicants Amendment filed on July 18, 2005.
Applicant amended *claims 2, 12, 16-18, 22, 29, 33-35, and 42-43*, canceled *claims 1, 11, 21, 28, and 36-39*, and added *claim 44*. *Claims 2-10, 12-20, 22-27, 29-35, and 40-44* are presented for further consideration and examination.

Claim Rejections - 35 USC § 102

2. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

3. Claims 40-44 are rejected under 35 U.S.C. 102(b) as being anticipated by Senator et al. (U5005761677).

4. With regard to claim 40, Senator discloses,

- *a non-volatile memory storing a first inode locating a first file in said storage; and*
(Senator, col.2, lines 13-46., col.4, line 45 - col.5, line 47)

Senator teaches of *"a single index node (inode) 301 having the inode number 39 is depicted; [and] an entry 313 in the file directory for this file system contains the inode*

Art Unit: 2145

number 39 pointing to inode 301 thereby associating the name 'NAME1' with the inode 301" (Senator, col.4, lines 53-57).

- *a block manager configured to copy said first inode to a second inode, wherein said block manager is configured to change said second inode in response to updates to said first file, and wherein said block manager is configured to atomically update said first file in response to a commit of said first file by writing said second inode to said non-volatile memory, whereby said second inode locates said first file in said storage. (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)*

Senator teaches of "a module, that in response to a system call argument to allocate another node in the file system tables and to copy the data block allocations from the old node into the newly allocated node. Both nodes now contain the same data block allocation information" (Senator, col.2, lines 15-19) referring to the same file.

According to Senator, "changes to the actual data are now made with respect to the new node" (Senator, col.2, lines 21-22); and that the newly allocated node corresponds to the updated file "after a COMMIT file operation or an FSYNC system call" (Senator, col.5, lines 6-7). Hence, "upon a read of [the file], it appears to the application program that the data in logical block has changed" (Senator, col.5, lines 36-38), since "the version module resets the inode pointer in the file directory entry 313 breaking the pointer to inode 301 numbered 39 and then sets the value in entry 313 to 40 pointing to inode 311 numbered 40" (Senator, col.5, lines 20-23).

5. With regard to claims 41-44, Senator discloses,

- *wherein the first file in the non-volatile memory is a first version of the first file, and wherein the block manager is configured to create a second version of the first file in*

Art Unit: 2145

response to a first write command of the plurality of write commands, and wherein the block manager is configured to atomically replace the first version with the second version in response to the commit command. (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)

- *wherein the first version is associated with a first inode, and wherein the second version is created by copying the first inode to a second inode and modifying the second inode, and wherein the atomic update is performed by writing the second inode. (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)*
- *wherein the storage is an object-based storage and wherein the plurality of write commands and the commit command address a first object corresponding to the first file. (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)*

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

7. Claims 2-10, 12-20, 22-27, and 29-35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Senator et al. (U5005761677) and in view of Zheng et al. (US006571259B1).

8. With regard to claims 2, 8, 12, 22, and 29, Senator discloses,

Art Unit: 2145

- *a non-volatile memory storing a first inode locating a first file in said storage; and*
(Senator, col.2, lines 13-46; col.4, line 45 -col.5, line 47)

Senator teaches of "a single index node (mode) 301 having the inode number 39 is depicted; [and] an entry 393 in the file directory for this file system contains the inode number 39 pointing to inode 301 thereby associating the name "NAMED" with the inode 301" (Senator, col.4, lines 53-57).

- *a block manager configured to copy said first inode to a second inode, wherein said block manager is configured to change said second inode in response to updates to said first file, and wherein said block manager is configured to atomically update said first file in response to a commit of said first file by writing said second inode to said non-volatile memory, whereby said second inode locates said first file in said storage. (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)*

Senator teaches of "a module, [that in response] to a system call argument to allocate another node in the file system tables and to copy the data block allocations from the old node into the newly allocated node. Both nodes now contain the same data block allocation information" (Senator, col.2, lines 15-19) referring to the same file. According to Senator, "changes to the actual data are now made with respect to the new node" (Senator, col.2, lines 21-22); and that the newly allocated node corresponds to the updated file "after a COMMIT file operation or an FSYNC system call" (Senator, col.5, lines 6-7). Hence, "upon a read of [the file], it appears to the application program that the data in logical block has changed" (Senator, col.5, lines 36-38); since "the version module resets the inode pointer in the tile directory entry 313 breaking the pointer to

Art Unit: 2145

inode 301 numbered 39 and then sets the value in entry 313 to 40 pointing to inode 311 numbered 40" (Senator, col.5, lines 20-23).

However, Senator does not explicitly disclose,

- *wherein said non-volatile memory stores a journal comprising a list of committed inodes, and wherein said block manager is configured to record said second inode in said journal.*

Senator teaches,

- *wherein said non-volatile memory stores a journal comprising a list of committed Modes, and wherein said block manager is configured to record said second inode in said journal. (Zheng, col. 1, line 66- col.2, line 8; col.6, lines 1-54)*
Zheng teaches of the use of a "list [that] is managed by inserting on the list a pointer to each cache block when the contents of the cache block are committed to the on-disk file system, which occurs in response to a commit request from the client" (Zheng, col.6, lines 28-31). Furthermore, according to Zheng, "the file system cache manager 34 also accesses an in-memory file system index 37, [which] includes indexing information for file system objects that have been added, deleted, or otherwise modified from the committed, on-disk state" (Zheng, col.6, lines 32-36).

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention was made to combine the teachings of Zheng with the teachings of Senator to provide an alternate method of "[recovering] from a system failure by restoring the database to its consistent state existing just after commitment of the last completed transaction ... [by maintaining] a log file of the database changes and the commit commands ... [including] a sufficient amount

of information (such as 'before' and 'after' images) in order to undo the changes made to the database since the last commit command" (Zheng, col. 1, line 67 - col.2, line 8). According to Senator, it is known in the art to recover from a system crash due to a power failure for example by *"[keeping] a file of records as they existed before [the changes] by the transaction, ... [or] to log each change in a journal or log file before the actual file system records are changed"* (Senator, col.2, lines 25-31).

9. With regard to claims 3-5, 18-20, 23-25, and 35, Senator and Zheng disclose,
- *wherein said commit of said first file comprises a commit command received from an external source which updates said first file.* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)
 - *wherein said commit command comprises a file close command.* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)
 - *wherein said commit command comprises an fsync command.* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)
10. With regard to claims 6-10 and 26-27, Senator and Zheng disclose,
- *wherein said journal further includes a checkpoint record including a description of an inode file, a block allocation bitmap, and an inode allocation bitmap* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47; Zheng, col.3, line 3 col.4, line 14; col. 14, line 46 - col.14, line 14)
 - *wherein the description comprises inodes for each of said inode tile, said block allocation bitmap, and said inode allocation bitmap.* (Senator, col.2, lines 13-46;

Art Unit: 2145

col.4, line 45 - col.5, line 47; Zheng, col.3, line 3 - col.4, line 14; col. 14, line 46
col. 14, line 14)

- *wherein said commit command comprises a rile close command* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)
- *wherein said commit command comprises an fsync command.* (Senator, col.2, lines 13-46; col.4, line 45 - col.5, line 47)

11. Claims 13-17 and 30-34 are rejected under 35 U.S.C. 103(a) as being obvious over Senator et al. (U5005761677) in view of Zheng et al. (US006571259B1) as applied to *claims 12 and 29* respectively above, and further in view of Raz (US005701480).

12. With regard to claims 13-17 and 30-34, Senator and Zheng disclose,

See *claims 12 and 29* rejections as detailed above.

However, Senator and Zheng do not explicitly disclose,

- *further comprising writing a master inode corresponding to an inode file including said second inode to a checkpoint record in said journal.*
- *wherein recovering from a system failure comprises:*
 - *scanning said journal to locate a most recent checkpoint record and zero or more inodes subsequent to said most recent checkpoint record within said journal; copying said master inode from said most recent checkpoint record to a volatile memory; and*
 - *updating an inode file corresponding to said master inode with said one or more inodes subsequent to said most recent checkpoint record.*
- *wherein said updating said inode file comprises:*

Art Unit: 2145

- *copying one or more blocks of said inode file storing said one or more inodes to a copied one or more blocks; and*
- *updating said master inode in said volatile memory to point to said copied one or more blocks.*

Raz teaches,

- *further comprising writing a master inode corresponding to an inode file including said second inode to a checkpoint record in said journal. (Raz, col.62, line 3 - col.63, line 44)*

Raz teaches of "an alternative 'redo' recovery mechanism, [where the] updated records are not flushed to state memory after every transaction. Instead, updated records are written sequentially to an after-image log, and all of the updated records are flushed to state memory only when certain 'check-points' occur. The check-points occur, for example, after a predetermined number of bytes have been written to the after-image log since the last checkpoint" (Raz, col.62, lines 18-26). Hence, "when a system crash occurs, the volatile state memory existing at the end of the last committed transaction is reconstructed by reading from the non-volatile memory the state memory records existing at the time of the last checkpoint, and redoing the modifications recorded in the after image log. The after-image log, for example, is read sequentially while redoing the modifications" (Raz, col.62, lines 28-34).

- *wherein recovering from a system failure comprises:*
 - *scanning said journal to locate a most recent checkpoint record and zero or more inodes subsequent to said most recent checkpoint record within said journal; (Raz, col.62, line 3 - col.63, line 44)*

Art Unit: 2145

- *copying said master inode from said most recent checkpoint record to a volatile memory; and (Raz, col.62, line 3, - col.63, line 44)*
- *updating an inode file corresponding to said master inode with said one or more inodes subsequent to said most recent checkpoint record. (Raz, col.62, line 3 - col.63, line 44)*
- *wherein said updating said inode file comprises:*
 - *copying one or more blocks of said inode file storing said one or more inodes to a copied one or more blocks; and (Raz, col.62, line 3 - col.63, line 44)*
 - *updating said master inode in said volatile memory to point to said copied one or more blocks. (Raz, col.62, line 3 - col.63, line 44)*

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention was made to combine the teachings of Raz with the teachings of Senator and Zheng to provide *"an alternative 'redo' recovery mechanism"* (Raz, col.62, line 18) in the event where a system crash occurs. According to Raz, *"an additional advantage is that the conventional state memory and snapshot memory caching facility can be used for maintaining the state memory cache and snapshot memory cache, and a conventional after image journaling facility can be used for maintaining the after-image log"* (Raz, col.62, lines 50-54).

Response to Arguments

13. Applicant's arguments filed July 18, 2005 have been considered but they are not persuasive.

Art Unit: 2145

14. The applicant argues in pages 12-15 generally for all independent *claims 2, 8, 12, 22, and 29* that the cited reference does not disclose a non-volatile memory for storing a first inode locating a first file and also for storing a journal comprising a list of committed inodes. The examiner respectfully submits that nowhere in either references disclose or mention the term "volatile memory" as argued by the applicant. Instead, the cited reference clearly defines the structure of possible to recover from a system failure by restoring the database to its consistent state existing just after commitment of the last complete transaction (e.g. col.1, line 65 - col. 2, line 8). This indication the memory must be non-volatile type of memory in order to store and recover from crash.
15. The applicant argues in pages 12-15 further for all independent *claims 2, 8, 12, 22, and 29* that the cited reference does not disclose the block manager is configured to record second inode in journal instead of providing various versions of a file without the need for data copy or log operations. The examiner respectfully submits that Senator does not need to make a copy of its original data, but it needs to keep a copy of its modified version of a file. Obviously, the modified version of a file is a copy of its original file and the differences.
16. The applicant argues in page 15 for *claims 8-10* that the cited reference does not disclose the features including: "*a storage coupled to receive said one or more write commands ... corresponding to said one or more write commands*". The examiner respectfully submits that the rejection is clearly stated the argued features. To reiterate, col. 2, lines 13-46; col. 4, line 45 - col. 5, line 47 clearly disclose the features including:

Art Unit: 2145

"a storage coupled to receive said one or more write commands ... corresponding to said one or more write commands".

17. The applicant argues in pages 16-17 for *claims 40-44* that the cited reference fails to disclose the limitations cited in claim 40. The examiner respectfully submits that the rejection clearly is made under U.S.C. 102 as seen above.

Conclusion


18. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a). A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.
19. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Thomas Duong whose telephone number is 571/272-3911. The examiner can normally be reached on M-F 7:30AM - 4:00PM. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Rupal D. Dharia can be reached on 571/272-3880. The fax phone numbers for the organization where

Art Unit: 2145

this application or proceeding is assigned are 571/273-8300 for regular communications and 571/273-8300 for After Final communications.

Thomas Duong (AU2145)

October 3, 2005


RUPAL DHARIA
SUPERVISORY PATENT EXAMINER